AI Reps & Warranties: Emerging Issues in Deals and Commercial Contracts

June 19, 2025
By Mark G. Wendaur, IV

Artificial intelligence quickly became embedded into business operations, software platforms, internal workflows, and consumer-facing applications. This means the legal risks associated with its development and growth are moving from the abstract to the real world. AI is not only changing how businesses operate, but also how leaders and legal practitioners must structure and negotiate contracts. Legal teams, in-house counsel, and M&A deal professionals can no longer consider this a niche issue and should consider it a negotiated deal point involving risk allocation, liability exposure, and asset valuation.

While there is significant attention paid to the disruptive operational power of AI, its implications on representations and warranties in commercial agreements and corporate transactions deserve just as much attention within the legal community.

Some tailored legal frameworks are already emerging to address the novel concerns around data privacy, intellectual property, indemnity, and operational continuity. These issues are especially critical in the context of software acquisitions, SaaS contracts, and any M&A deal involving AI-derived intellectual property or business processes.

AI-Specific Representations in Deals

The National Venture Capital Association (NVCA) model forms were updated at the end of 2024 to incorporate AI-specific representations and warranties. These terms are increasingly reflected in market practice:

- Targets must affirm that AI tools were used in compliance with applicable licenses, regulations, and data use agreements.
- Targets must represent that they did not input any personal, confidential, or protected information into AI tools, unless those tools guarantee that such data is not used for training or product enhancement.
- Targets must represent that data was deidentified or anonymized prior to use in training models to avoid running afoul of applicable data protection standards.
- Targets must disclose any generative AI platforms used to develop proprietary IP and warrant that such use does not jeopardize any ownership rights being acquired.

The above issues are only the tip of the iceberg. Transactions and agreements involving complex AI products must include increasingly technical reps concerning:

- Model training logs and documentation
- Fine-tuning methods and retention of model weights

- Mechanisms used in retrieval-augmented generation (RAG)
- Use of synthetic or auto-generated content in commercial workflows
- Model validation performance thresholds, including floating point operation limits and accuracy ranges

As the technology matures, the legal community is catching up by embedding operational guardrails directly into transactional documentation.

Practical Contractual Risk Areas in AI Licensing

Commercial licensing agreements involving AI must now be viewed through a much more detailed legal lens. Counsel should be prepared to negotiate contract terms specifically addressing:

- Non-infringement guarantees concerning both source code and training data, especially where data scraping or aggregation may have occurred
- Explicit ownership claims over AI outputs, derivatives, and model weights
- Training data auditability, including the legal basis for data collection, classification, and labeling
- Compliance with global privacy and cybersecurity laws, including GDPR, CPRA, and HIPAA when applicable
- Robust indemnification obligations for breaches of data usage restrictions, IP violations, or algorithmic harms
- Tech E&O and cyber liability insurance provisions, ensuring recourse exists if generative tools malfunction, hallucinate, or produce defamatory content

In many cases, the liability profile of AI is uncertain and difficult to quantify. Because many models operate as "black boxes," licensees are often left without a clear explanation of how certain outputs were generated or what datasets were used in training. This makes traditional warranties about performance or fitness for a particular purpose difficult to enforce.

As a result, buyers and licensees are demanding broader representations, heightened disclosure obligations, and post-closing audit rights to mitigate these unknowns (while sellers are seeking to disclaim warranties and narrow representations).

M&A and AI Due Diligence

AI risk has quickly become a key diligence category in M&A deals, especially in transactions involving software, e-commerce, analytics, or consumer engagement platforms. Buyers are now expected to conduct diligence not just on IP rights and customer contracts, but also on how AI has been implemented and governed.

Pre-Acquisition Diligence

Key diligence areas include:

- Training data sourcing: Was the data obtained lawfully and under enforceable terms?
- Privacy risk: Was any personally identifiable information (PII) used in training or prompting without consent?
- Third-party code and APIs: Does the AI product depend on third-party components that might limit assignability or trigger license fees?
- Model update and retraining rights: Who controls the model lifecycle, including patches and performance tuning?
- Export control risks: Could the AI model be subject to ITAR, EAR, or other national security controls due to its capabilities?

Post-Closing Continuity

Buyers should also require:

- Complete AI architecture diagrams and component inventories
- Documentation of ethical safeguards and bias mitigation processes
- Retention policies around input prompts and AI-generated output logs
- Model deployment playbooks and downtime risk disclosures

Transition services agreements, software escrow, and founder retention may also be needed to ensure business continuity and proper knowledge transfer where AI is a critical but complex asset.

IP, Privacy, and Employment Triggers

This isn't only an issue for "tech transactions." As AI expands across business functions, its legal implications multiply. Core issues include:

- Intellectual property ownership, particularly whether AI-generated outputs are protectable under U.S. copyright law or must be secured as trade secrets
- Privacy and cybersecurity risks stemming from unstructured data ingestion and prompt leakage, especially when sensitive information is processed or used
- Employment law exposure, including discriminatory hiring algorithms or opaque automated decision-making processes that may violate EEOC or state-level labor rules

Recent case law and regulatory action suggest that companies using AI for decision-making will be held to explainability and fairness standards, even if they do not fully control or understand the model.

Additionally, companies leveraging AI in consumer products or safety-critical environments must consider product liability exposure under traditional tort theories, especially if AI contributes to physical or economic harm. Think about the auto-driving taxi that must decide whether to hit the pedestrian or crash the car.

AI and IP Security Agreements

As more companies develop proprietary AI tools, models, and datasets, lenders and investors are increasingly taking security interests in these intangible assets. This requires a rethinking of traditional IP Security Agreements.

When collateral includes AI-generated or AI-driven intellectual property, legal teams should evaluate:

- Whether model weights, training datasets, or prompt libraries are clearly documented and listed as pledged assets
- Whether the borrower can demonstrate ownership and provenance of the training data and model code
- If the model is fine-tuned from a third-party foundation model, whether the underlying license permits encumbrance or assignment
- If retrieval-augmented generation (RAG) is used, whether the underlying corpuses and connectors are part of the security package
- The existence of source code escrow to ensure access in the event of default or bankruptcy
- Any restrictions in open source or SaaS agreements that may limit foreclosure or reassignment rights

Moreover, the lender’s enforcement rights may be limited if the AI model or data is co-owned, cloud-hosted, or reliant on third-party APIs. Security interests must be carefully drafted to reflect operational dependencies, and perfection of those interests

may require filings beyond the USPTO, including notice to cloud vendors or consent from licensors.

IP Security Agreements for AI assets must go beyond standard boilerplate and should be tailored to the unique hybrid nature of AI systems combining software, services, and data streams. In many cases, a supplemental AI-specific collateral schedule may be appropriate.

Takeaways for Legal and Deal Teams

AI is no longer a novel technology element to be glossed over in standard reps and warranties. It is a high-stakes business driver that intersects with every major legal category: IP, privacy, cybersecurity, employment, antitrust, and contract liability.

Actionable takeaways include:

- Draft AI-specific reps and warranties that cover data sourcing, training protocols, model rights, and use case restrictions
- Build diligence frameworks that include discussions with technical teams and review of logs, policies, and product roadmaps
- Negotiate indemnification mechanisms that allocate financial risk from misuse, error, or regulatory exposure
- Ensure insurance provisions cover AI incidents, from hallucinated content to data leakage
- Establish post-closing governance and monitoring structures, particularly in acquisitions involving live AI models or mission-critical algorithms
- Review and update IP Security Agreements to specifically address AI collateral and embedded third-party dependencies

Ultimately, the central legal question becomes: When AI makes a mistake, who pays? Whether drafting a commercial SaaS agreement or executing a strategic acquisition, every deal team must be ready to answer that.

As legal and technological standards continue to evolve, ongoing adaptation will be essential.